

# USING OBJECT-ORIENTED ANALYSIS TO DESIGN A MULTI-MISSION GROUND DATA SYSTEM

Peter Shames  
Jet Propulsion Laboratory  
Multi-mission Image Processing Laboratory  
Pasadena, California

## **Abstract**

This paper describes an analytical approach and descriptive methodology that is adapted from Object-Oriented Analysis (OOA) techniques. The technique is described and then used to communicate key issues of system logical architecture. The essence of the approach is to limit the analysis to only service objects, with the idea of providing a direct mapping from the design to a client-server implementation. Key perspectives on the system, such as user interaction, data flow and management, service interfaces, hardware configuration, and system and data integrity are covered. A significant advantage of this service-oriented approach is that it permits mapping all of these different perspectives on the system onto a single common substrate. This services substrate is readily represented diagrammatically, thus making details of the overall design much more accessible.

## **1.0 Context**

The Multi-mission Image Processing System (MIPS) performs science instrument data processing and archival product generation for high data rate instruments flown on a wide variety of planetary missions. The system now supports the Galileo, Cassini, and Mars Pathfinder missions and is planned for support of a number of Discovery-class missions. It is part of the Advanced Multi-Mission Operations System (AMMOS) developed at the Jet Propulsion Lab (JPL).

The design integrates a variety of vendor platforms, commercial sub-systems, and customized components into an evolvable system capable of handling a broad range of spacecraft data rates and processing requirements. Support is provided for both local and remote users via workstations, and for remote users without local processing capabilities. Major parts of the system are

portable and re-configurable, anti a substantial legacy of image processing, analysis, and display applications, called VICAR, are being ported as part of the development task.

The system has just been redesigned from a centralized, VAX-based architecture to a Unix-based, fully distributed, client-server architecture. Part of the intent of the redesign was to move the system into these more modern technologies and to make the new system more easily adapted to new mission processing profiles. Modeling, benchmarking, and testbedding have been extensively used to demonstrate feasibility and the design has successfully passed Critical Review.

## **2.0. Motivation**

At the start of the MIPS re-design project there was no overall description of the existing architecture of the system nor of how existing elements, and the new ones that were contemplated, would integrate into the final design. There was, however, a substantial legacy of existing code and sub-system components to be integrated that had been part of the 2nd generation VAX system. While most of the technical and operations staff were familiar with VAX systems, only a few understood Unix and distributed client-server systems.

The overall system design had to be accessible from several perspectives, starting with the user/operator view and including logical, physical, and functional ones. Hardware, software, user services, networks, security, and access control elements all needed to be made visible and presented in context. Existing components and processes had to be shown within the context of the new architecture and interfaces. Accordingly, a major part of the system design effort went into defining the

overall architecture, communicating that design to the senior user and technical staff, and then refining the design with their active involvement.

The initial approach used high-level dataflow diagrams and context diagrams to describe the overall system design. This was useful from an analytical perspective, but did not convey the design with enough clarity to those who were unfamiliar with client-server architectures. Following some research into object-oriented analysis (OOA) and design techniques the present approach was developed. This approach was strongly influenced by the techniques defined by Coad and Yourdon in their book Object Oriented Analysis, but uses a service-oriented perspective instead of a pure object-oriented one, and adds analytical framework and descriptive and representational approaches that have pragmatic applicability for a broad class of re-design efforts.

This paper describes the adaptation of this Object Oriented Analysis (OOA) technique and then demonstrates using it to communicate key issues of system logical architecture. Key perspectives on the system, such as user interaction, data flow and management, client-server interfaces, hardware configuration, and system and data integrity are covered. A significant advantage of this service-oriented approach is that it permits mapping all of these different perspectives on the system onto a single common substrate. This services substrate is readily represented diagrammatically, thus making details of the overall design much more accessible.

### **3.0 Genesis of the Technique**

The OOA technique presented by Coad and Yourdon is applicable to a complete system analysis task where a basic design is being newly developed, but where there may be some information derived from previous designs. They describe an analytical process that involves identification of objects, their structure, the subjects, their attributes, and finally the services performed by the system. In their approach objects are understood to encapsulate attributes and services on those attributes;

objects with message passing interfaces are the basic entities that appear in the design abstraction.

Coad and Yourdon describe objects as the least volatile elements of the design and posit them as a stable framework upon which to base the rest of the analysis and specification strategy. Other items in this framework, data, sequencing of functions, the functions themselves, and the interfaces among system components are seen as having increasingly higher levels of volatility. Object definitions are proposed as the basis of the analysis, and these definitions include a specification of the services provided by each object.

Interfaces have a high volatility in this view because they are implemented as messages passed between objects. The message passing functions are a stable part of the design, but the message formats are bound late in the analysis and remain fluid.

The process of identifying objects involves looking for logical structures, other external systems, devices, events that must be remembered, roles played, locations, and organizational units. These elements are each reviewed to identify common services, attributes, remembered information, and requirements. Services, in their technique, are specified last because they are seen as being highly volatile, including not just processes but also the details of process sequencing. A service is simply defined as "the processing to be performed upon receipt of a message".

Services are identified during system analysis by considering the behavior of objects and the basic functions each provides. Behaviors include:

- Immediate causation - state, event, response
- History - object life history
- Function - fundamental services

The basic services that an object provides may include:

- Occur (instance add, change, delete, or select)

Calculate (create, process, alter, derive)  
Monitor (observe, record, track)

Where a new system is being designed from the ground up, or an old system is being completely redesigned, such a top-down object-oriented analysis is likely to be extremely useful, especially if the implementation is to be done in an object-oriented language such as C++ or ADA. It is less clear how a complete object-oriented analysis will benefit a more modest re-implementation or system transition project such as the one that prompted this activity. The resulting OO design may be highly descriptive of the problem, but mapping this design onto existing services and legacy applications is likely to be problematic.

The heart of the modified approach is to adopt a client-server architecture and to focus the analysis on only one class of object, Services. A client-server architecture maps cleanly onto the services perspective, providing a useful analytical and descriptive framework.

### 3.1 Client-Server Architectures

Client-server implementations are now almost commonplace and many elements in modern networked systems use this technology. Many capabilities provided in distributed computing environments, such as networked services, database servers, file systems, and information resources, all incorporate some form of client-server technology. Common services like the TCP/IP remote login (telnet) and file transfer protocol (FTP), the Sun Microsystems developed Network File System (NFS) service, and commercial relational DBMS servers like those from Sybase and Oracle all use client-server approaches. A major component in the Open Software Foundation (OSF) system is the Distributed Computing Environment (DCE), a framework for developing client-server applications,

client-server approaches are "in" at the practical level of developing and integrating systems based upon legacy applications, commercial server offerings, and current technology, but there are no well accepted

approaches for defining and describing such systems. The older design technologies, such as those which were useful for more centralized implementations, fail to capture much of the richness of a client-server approach or to provide an integrated description of the system. Key elements here are:

- Distributed nature of the system
- Encapsulation of functions and data
- Stable and clearly specified interfaces between components
- Reuse of services and interfaces in the design

Many of these elements are considered in an object-oriented approach, the key difference here being a focus upon services as a basic substrate, rather than all other classes of objects. From this perspective services are relatively stable entities, tied to fixed interfaces and protocols, existing applications and/or capabilities, and commercial servers. The volatility and adaptability in such a system tends to be embodied in new processing algorithms and in the scripts that drive the system, rather than in the definition of the services themselves.

Particularly in our systems, which are used for science data analysis as well as production processing, we assume that the scripts will change from moment to moment (during interactive analysis) or that they may persist for the duration of some project (production data processing). Most of the services are stable over time, though specific details of database content or even data type will change. Where processing details must also change the services framework provides the means for accommodating new user interfaces, algorithms, and data formats. Much of the processing for new projects or new products can be accommodated by adapting existing components and scripting or adapting the processing scenario.

### 3.2 Basic Service Element

In this approach the essential building block is a basic service element (BSE). Figure 1 provides a diagrammatic view of this model. Each service element consists

of the software, hardware, and staff that are required to carry out that function. Each service element has a defined protocol interface and three primary protocol elements: request, response, and fulfillment. In most cases the protocol has an electronic implementation with appropriate protocol data units, but the model does not preclude such options as a written request and a written response or a verbal request and physical media shipped in fulfillment.

and] a file defining the processing to be carried out and the account to be charged.

Response - the reply or answer to the request, stated in an electronic message or other medium of exchange. For a film recorder service the response consists of a message that either accepts the request or denies it for some reason, such as inadequate funds or privileges.

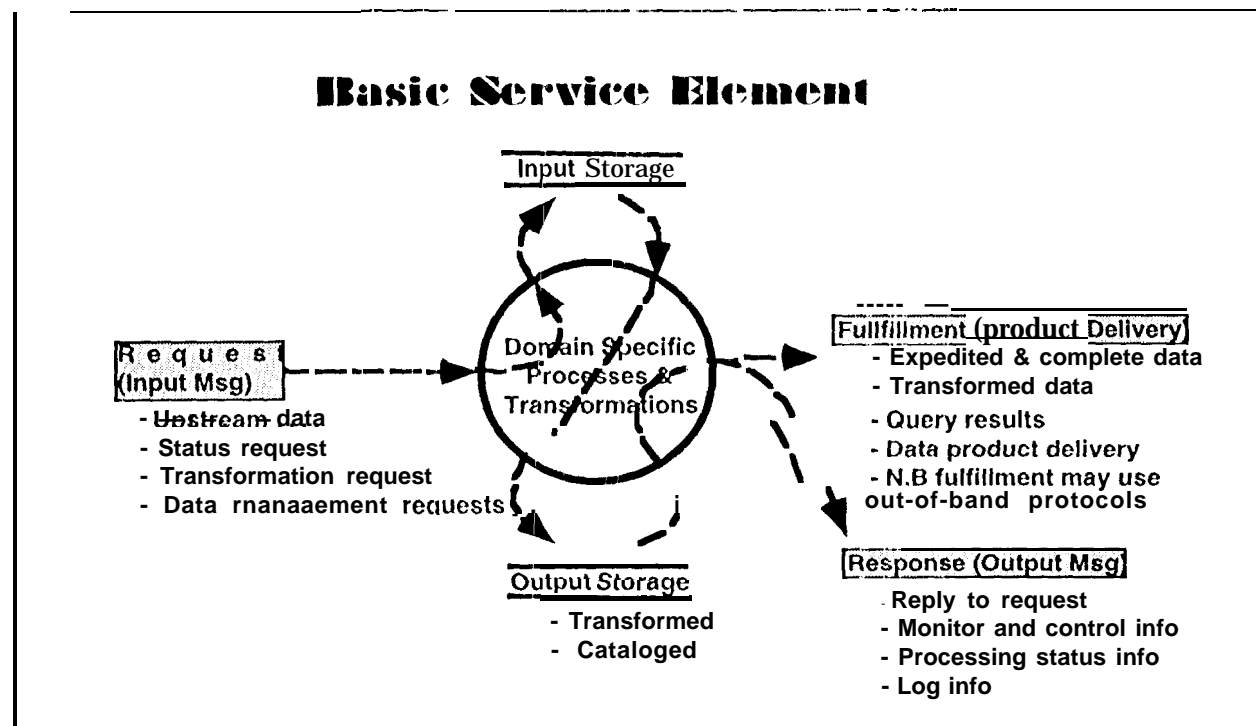


Figure 1

The components of a basic service element are defined as follows:

Protocol - a definition of the process by which requests are made and response and/or fulfillment of the request is provided.

Protocol Data Units - the actual items that are passed during a protocol exchange. May be defined electronic message formats, physical media, or other items of exchange.

Request - the service that is being asked for, stated in an electronic message or other medium of exchange. For a film recorder service the request consists of a data file

Fulfillment - the completion of the request or delivery of a product, often in some other medium of exchange than electronic. For a film recorder service the fulfillment consists of delivery, via mail, of the requested film products.

In many cases the fulfillment of the request may be combined with the response. An example of this might be the return of a the requested rows from a table in response to a database query.

Each service element provides request-driven delivery of data, supports multiple concurrent clients, and uses simple data driven queries. Each element provides its

own transform functions, local buffer or cache, and implementation approach, based upon requirements. The Internal implementation choices are opaque to the external user and two implementations of the same service, meeting the same interface protocols, may be entirely different internally.

### 3.3 Analyzing System Services

The analysis process begins by identifying the basic processes and transformations that the system embodies. This includes an analysis of the essential structure of the processes and the data, identifying major data flows, key transformations, and essential products. As shown earlier, the basic services that a system provides will typically include instances of Occur, Calculate, or Monitor functions.

The analysis should identify both services that are essential for performing the work of the system (creating products of whatever sort) and services that are essential for supporting the system (monitoring, tracking, and sustaining). It may prove useful as part of the analysis to group like services, based upon related processes, the class of data handled (image, textual, tabular), the type of service rendered (transformation, storage, monitor), or other characteristics such as temporal relationships or spatial location.

### 3.4 Developing the Services Substrate

The analysis of the system into services with well-specified interfaces is the real meat of this technique, but being able to communicate the resulting design, and its several perspective views, is an essential part of the strategy. The intent here is to make what is a multi-dimensional system structure accessible to others, some amount of experimentation with visual representations is likely to be necessary before an arrangement is arrived at that adequately communicates the design.

The process is to create a set of named icons representing the services, or service groups, and then to arrange them in some visually meaningful way. One approach is to organize them by major data flows through

the system, so that some sense of process flow is inherent in the arrangement. Another possible approach is to organize groups of services according to classes of processing (interactive, batch, data driven) or even according to organizational element, if that happens to make sense in the application context. There has not yet been any useful heuristic identified for facilitating this part of the process.

The end result of this effort should be a visual representation of the essential structure of the system, displayed as a SCt of services or service groups. Each of the entities on this diagram represents one or more instances of a basic service element, with specific processing, local data storage, and associated interface protocols, and protocol data units. This services perspective should show all of the major components of the system, labeled as to the function each provides. This basic services perspective may then be treated as a substrate on which to map many different perspectives on the system.

Figure 2 shows the result of applying this technique to a high-level analysis of the AMMOS spacecraft ground data system. Clearly each of the "blobs" in this diagram represent more than just one service. In fact they represent groups of associated services, where each group has a related set of functions and operates on a common set of data. Such a high level diagram is useful for providing context for the whole system; lower level decomposition into individual services and interface permits the necessary levels of detail and specificity to be exposed.

### 3.5 Mapping Other Perspectives

Given the basic services substrate a number of other perspectives on the system can be presented. One key perspective on any major data handling system are its data flows. One or more overlays, showing major data flows during different processing stages, at different time intervals, or in different operating modes can be used to explain the operational characteristics of the system. Different processing stages (telemetry handling, assembling level 0 instrument frames, production of CDR-ROMs with calibrated data), different processing

modes (batch, interactive, data-driven), or different user perspectives (remote science, production processing, operations) may require quite different representations or perspectives.

choices. While this approach provides a useful means of tracking and visualizing alternatives it does not obviate the need for careful system engineering studies, hardware sizing, or configuration analysis.

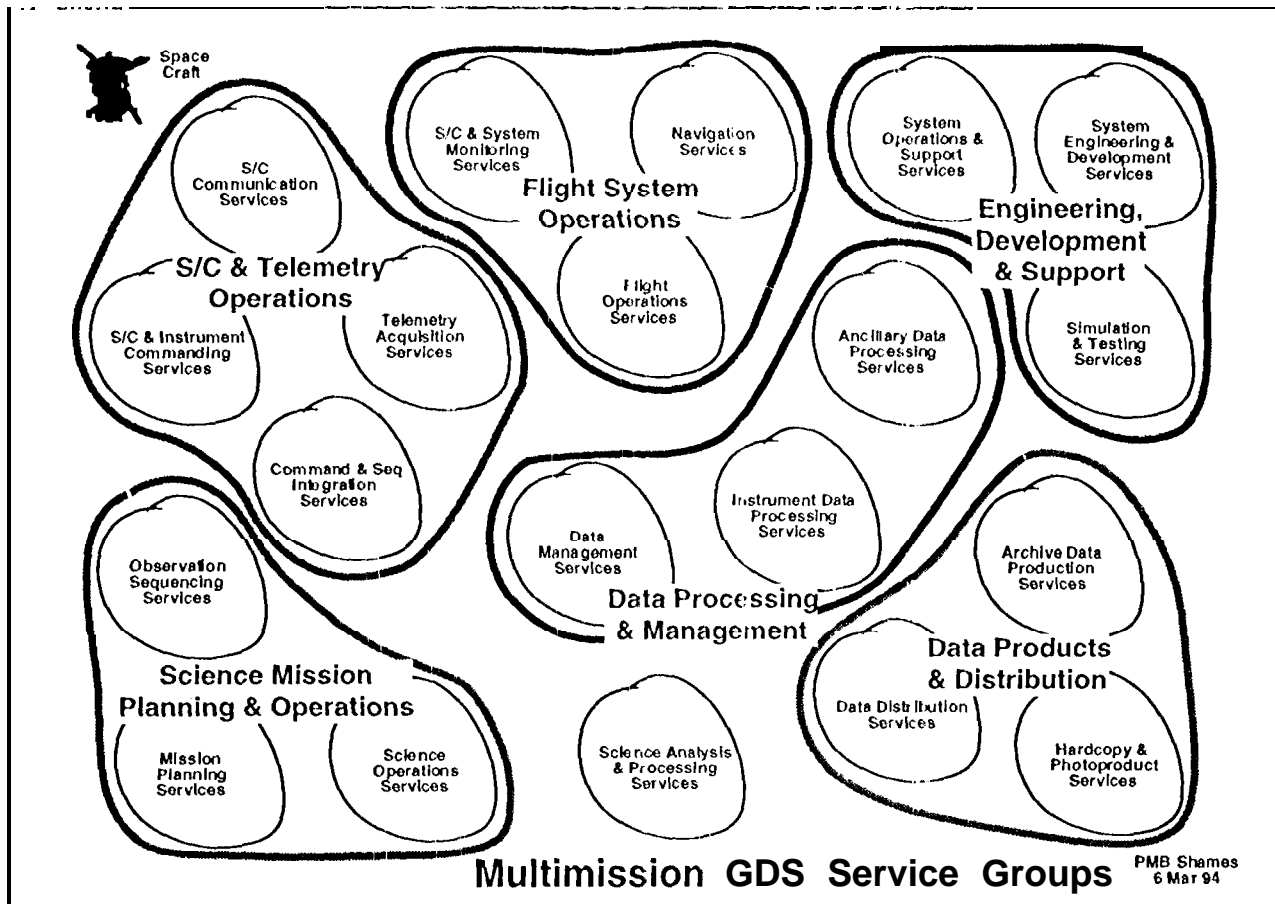


Figure 2

The software components that support each service can also be directly mapped onto the basic services substrate, thus showing this perspective of the system. Because user interfaces, service APIs, and other components may be used in several parts of the total system, re-use of S/W components should be readily evident in this view.

Mapping the system services onto hardware platforms may also be readily demonstrated, thus presenting the other major development aspect of an overall design. It is convenient to experiment with mapping the services design onto different hardware platforms, permitting an easy graphical assessment of different design

The basic substrate may also be used to analyze or represent other elements of an overall process design, such as organizational structures. It has proven instructive to map the different organizational entities involved in development, integration, and operation onto the basic service diagram. This can help clarify a variety of issues with regard to organizational boundaries and interfaces, clearly a part of the analysis of system and processes.

We have experimented with using this basic approach to elucidate other aspects of the overall system, with mixed results. For example, showing the data flows between

each service and a centralized monitoring subsystem produces a diagram that looks" like the proverbial plate of spaghetti.

Physical interconnections, such as local or wide-area networks, can also be overlaid on the basic substrate. In practice this has not proven to be too practical, largely for topological reasons related to security and server/network isolation considerations. Attempting to display control domains, particularly where there are complex overlapping access policies, with different styles and modes of access, has also not been too successful.

### 3.6 Specifying Interfaces and Protocols

The approach just described produces an accessible pictorial representation of the system design, one that is quite readily understood by a broad range of individuals. To complete the system description at a structural level it is also necessary to specify the interfaces, protocols, and protocol data units for each of these service entities. We have used two related organizing approaches for this, one is a simple template, adapted from the, ISO protocol stack, that describes the interface layers for each service, the other is a summary table of services, protocols, and protocol data units. An instance of this template, or mini-System Interface Specification (SIS), is created for each service defined in the design.

The mini-SIS departs from the 7-layer ISO stack in that it differentiates the application layer from an API or library call layer, and it collapses the transport and network layers. The modified stack has the following layers:

#### **Application/Utility**

Application Programs: These are programs that are developed as part of the interface, using lower level services and interfaces. Utilities may be provided as part of the interface, either as examples or templates, monitor and edit tools, or as primitive applications that exercise the API.

#### **API/Library**

Application Programming Interface (API)/Library: Developed in house or

provided by some other means. There may be one or more of these interfaces that are used at different levels in applications (OS interfaces, vendor interfaces, encapsulation interfaces).

#### **Presentation**

Data Representation: fundamental data format and representation (binary data formats, VAX, IEEE), also may reference data conversion interface layers such as External Data Representation (XDR).

#### **Session**

Authentication and Session Control: establishing connections, identifying communicating parties at each end, managing the session, handling of users, permissions, and access control. This could involve passwords, other means of electronically identifying individuals, or a handshake with a live user.

#### **Transport/Network**

Reliable End-to-End Data Transfer: means of moving data reliably from initial source to final destination. For networks these are the protocol layers such as the Transmission Control Protocol (TCP) and the Internet Protocol (IP) from TCP/IP. Where physical media such as CD-ROMS are being transferred this might be UI'S.

#### **Data Link**

Point-to-Point Link: protocol and link layer that provides a point to point connection and control. For networks this is the protocol layer such as Ethernet interface and encoding standards. For physical media this is N/A.

#### **Physical**

Actual Physical Medium: the medium which is being used to complete the interface. For networks this is the Ethernet coax, twisted pair, or optical fiber. Where actual physical media are being transferred this might be a CD-ROM or magnetic tape.

The other supporting document is a table describing each server (hardware platform), the services allocated to it, the access controls imposed for each service, and any related security concerns. Separate diagrams may also be provided showing the physical hardware configuration, both

overall and in detail, the network topology, and the physical layout of the facility. These diagrams are directly related to the hardware overlay perspective on the system and provide additional details on the servers themselves.

Some other perspectives on the system, such as network topology, isolation, access controls and other security concerns, appear to be most readily presented on a physical view of the system rather than a services one. By using common nomenclature these views may be directly related to the coherent view of the system provided by the services analysis, thus tying the entire design description together.

developed during the MIPS re-design. The toJ-level analysis of the MIPS system identified three major groups of processing services: telemetry and data processing, data and information management, and data product generation. It also identified two groups of support services: operations and management, and system support, and a set of services associated with users, either local or remote.

Figure 3 shows these service groupings and the individual services contained in each. Note that system engineering, development, and related support services are an essential part of MIPS but are not shown.

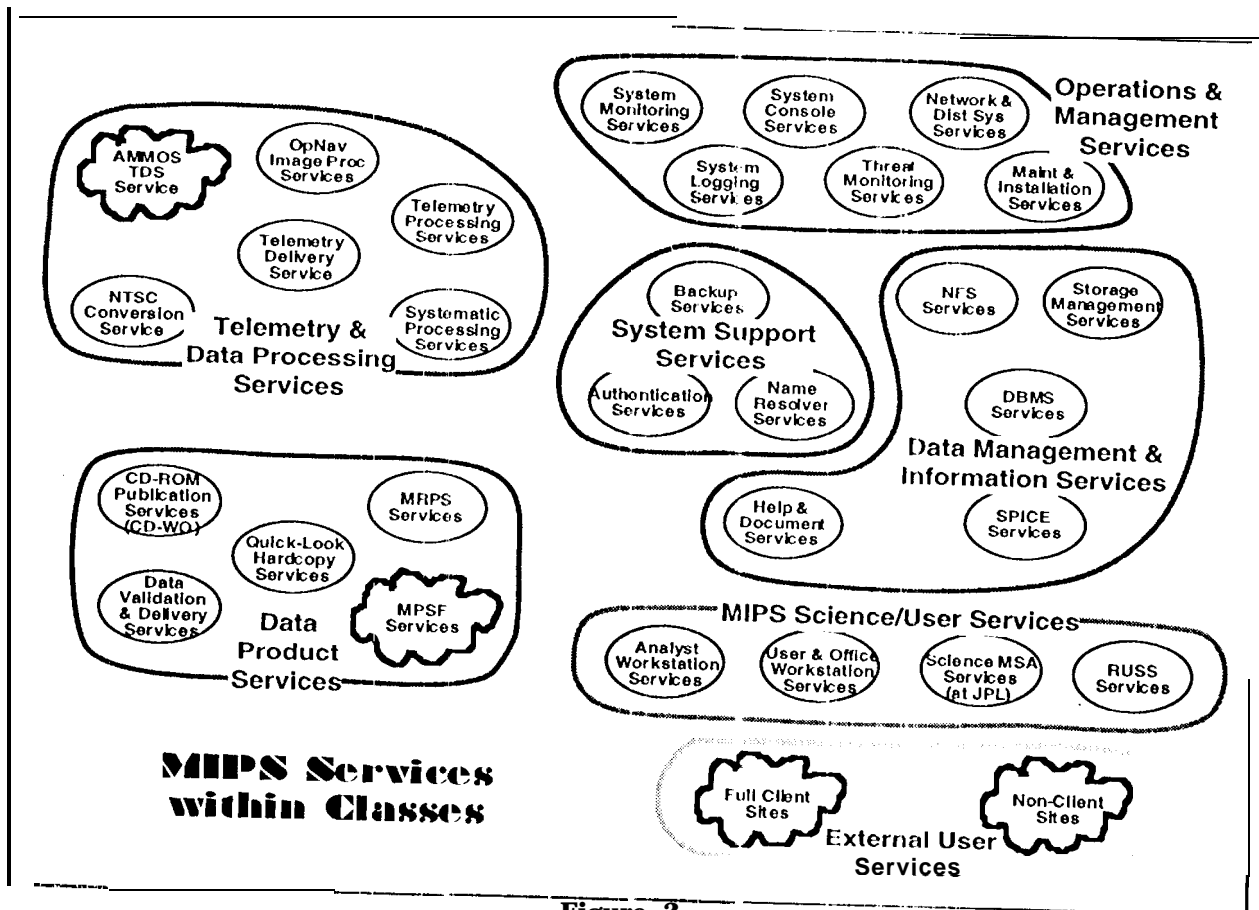


Figure 3

#### 4.0 Applying the Approach to a Real Problem

This section will demonstrate the application of this approach to a real design problem, using some of the materials

The services are grouped together based upon natural affinities, as should be obvious from the group names. There are of order 5-7 groups each with roughly 5-7 services defined. This "fingers of one hand"



or "span of cent ml" rule is commonplace and requires little further discussion. Do note that applying the technique to larger systems is possible, but that higher level aggregations of service groups may be required (as shown in Figure 2).

earlier, the effect of different allocations is easy to visualize using this representation, a capability we made use of during design trade studies.

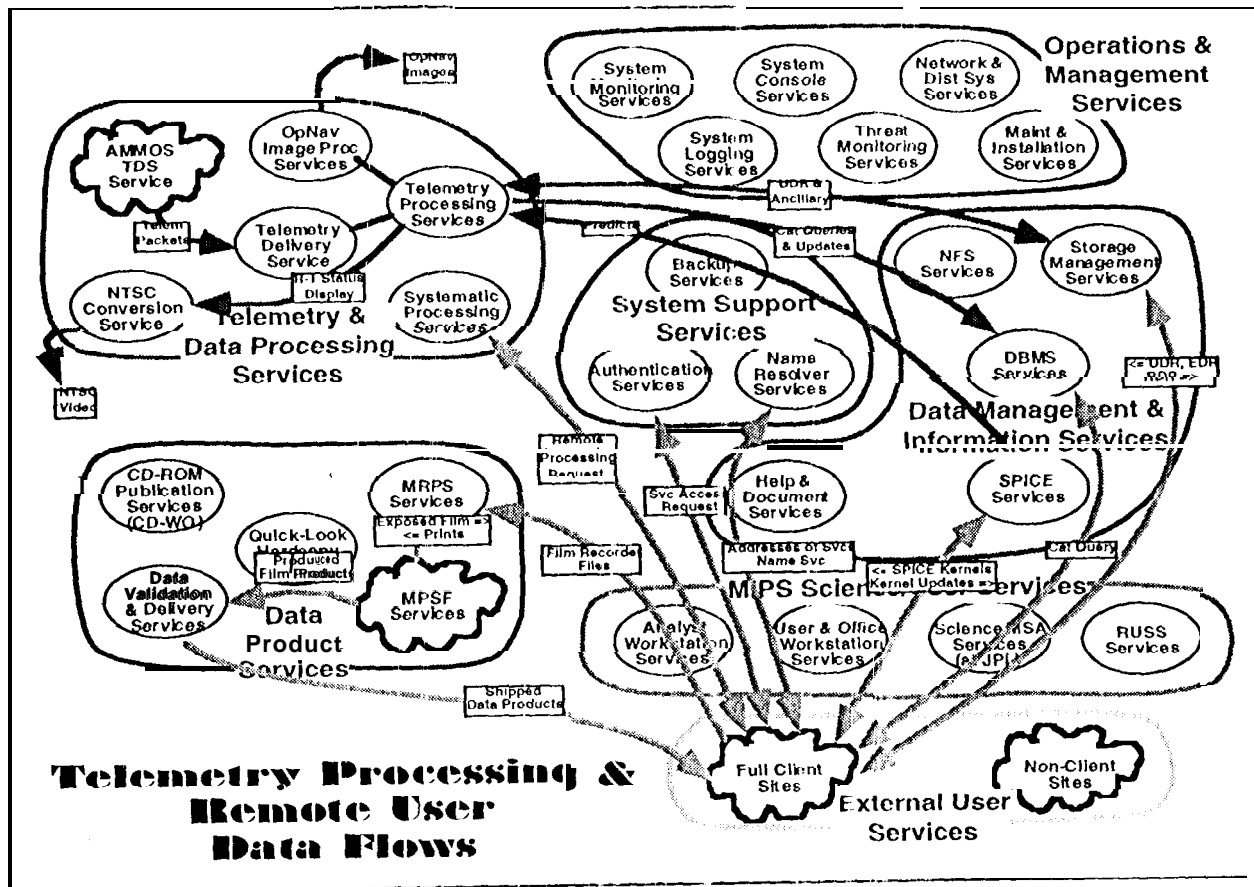


Figure 4

#### 4.1 Mapping Dataflow, Software, and Hardware

Figure 4 shows two different dataflow overlays on the basic MIPS service substrate. The upper dark arrows in Figure 4 demonstrate telemetry processing flows during real-time processing, while the lower light arrows show dataflows to a remote user at a workstation. Note that the arcs are labeled as to the key items that flow and that directional arrowheads provide additional information.

Figure 5 shows the mapping of services onto the physical hardware platforms that were used to implement the system. As noted

Figure 6 shows the mapping of services onto the software components that implement the system. Note that there is considerable re-use of software components and this perspective makes these relationships clear.

#### 4.2 Mapping Remaining Perspectives

While it is possible to use this technique to describe other perspectives of the system, not all of them have been as successful as the previous examples. As described in section 3.5, we have tried to use this approach to provide representations of monitoring subsystems, network topology, and access control domains, among others. These diagrams have not proven to be

particularly useful when mapped onto the services substrate, so we are not including them.

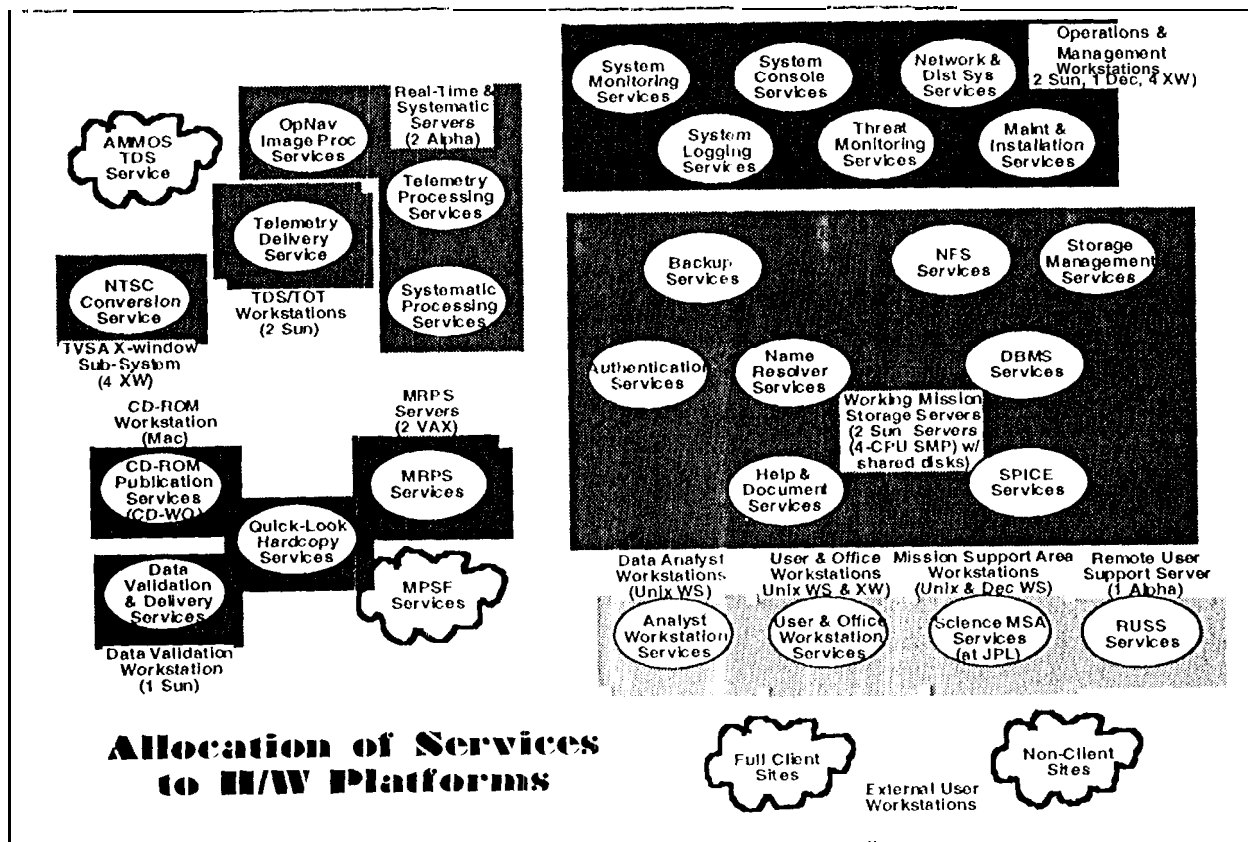


Figure 5

For the hardware related perspectives of the design a physical representation, which may be directly related to the hardware overlay, provides a better perspective. Both network topology and physical access control, which are implemented, in large part, at the physical level are best shown on a physical substrate. We have also not produced any represent at ion which is particularly useful for describing one to many relationships, such as those between a monitoring subsystem and each other service.

#### 4.3 Specifying the Service Interfaces

Table 1 shows the complete description of the service interfaces and protocols. Only some of the services are shown in this abbreviated view.

Table 2 shows the access control domain information used to specify in detail which physical servers run which services, who may access these services, and other details about how the security approach is defined and implemented.

A one page mini-SIS that uses the modified ISO template to define all interface layers is developed for each of the service interfaces.

#### 4.4 Adapting the design to new missions

The overall MITS system has been designed for adaptability and this latest re-design effort has further enhanced that by defining a new client-server architecture, with cleanly established interfaces. Starting from this baseline it has been possible to adapt the system to new missions and to new processing scenarios as required. The

architecture accommodates re-use of components and whole sub-systems, depending on the requirements of the new mission.

Most of the work of describing the system adaptation requires changes in only one or two overlays, such as data flows or

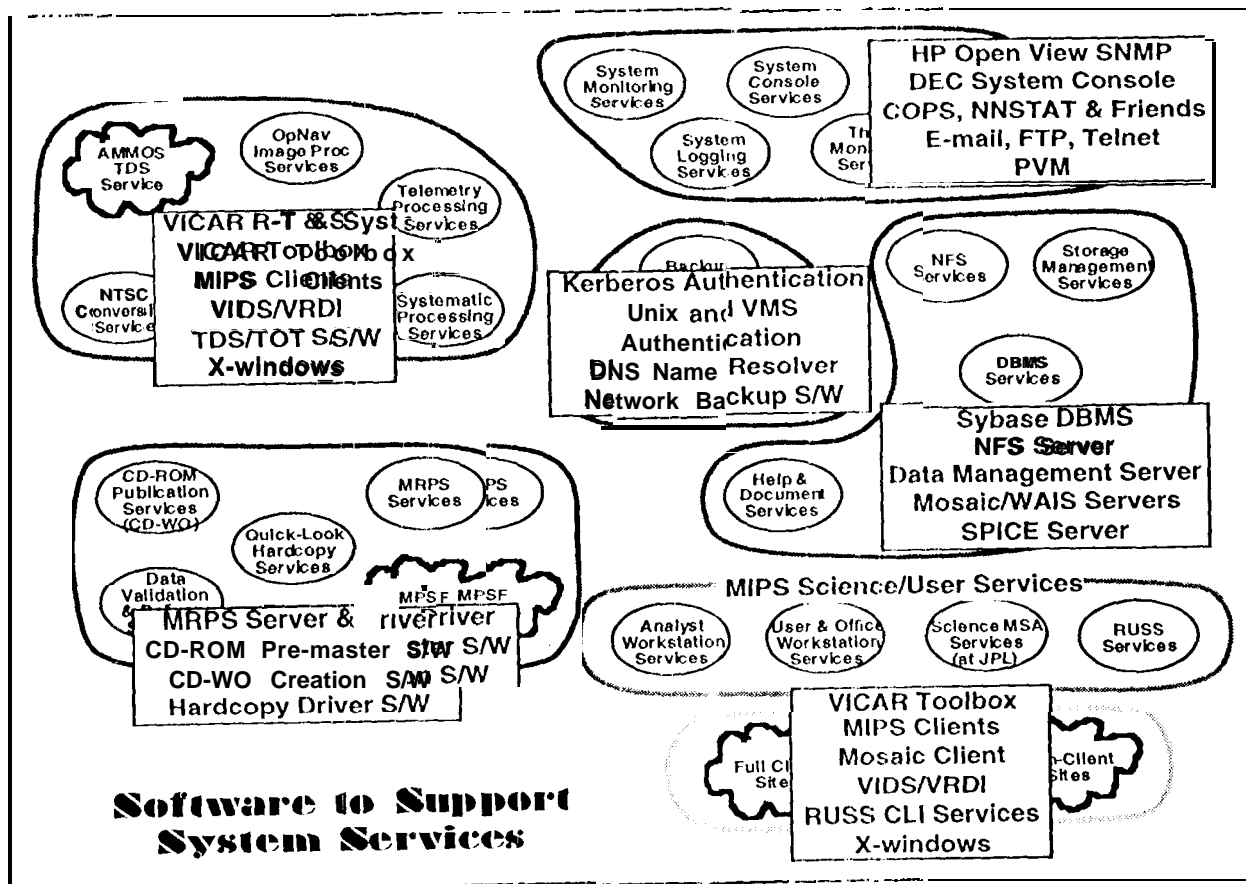


Figure 6

For some new missions, such as Mars Pathfinder with its very low data flows, most of the existing subsystems are used directly. The main telemetry and systematic processing is allocated to its own workstation, but the rest of the services are just adapted to the new data types. For other high rate missions, such as Cassini, some of the hardware platforms will need to be upgraded and larger data caches will be required. None of this affects the basic architecture of the system, though some specific processing algorithms, and many of the scripts, will be re-written to meet the new mission requirements.

Since the service-based design is largely invariant over such re-uses it is easy to use it to manage and describe these adaptations.

hardware platform allocations. The basic system processes, architecture, interfaces and protocols (at several levels), software design, and security structures do not change.

In the actual implementation, or during design trades, it is possible to update an individual service without altering the overall architecture. Because each service has defined interfaces at each of level in the protocol stack (data units, API, representation) it is possible to replace one instance of a service with another. A familiar example would be replacement of one relational DBMS server, which uses SQL and the standard API, with another one that adheres to the same interfaces but has

higher performance or improved failover characteristics.

This paper does not deal with many implementation details, but it should be noted that the definition of service interfaces, APIs, and protocols depends heavily upon open systems standards. Our standards include specification of Internet RFCs, POSIX, CCSDS, and Other non-proprietary national and international standards. Over time we expect the client-server design, this design representation, and the use of standards for implementation to pay a high return on investment.

### 5.0 Conclusion

As has been demonstrated, this design approach works very well for describing the key features of a complicated, distributed system. The technique has been shown to work well for the major MIPS system redesign and it has also been applied to the system analysis and re-design of the much larger AMMOS ground data system. Once the design has been described and the interfaces bound, moving into the implementation phase has been quite straightforward. Because this analytical approach and the associated descriptive technique makes the features of the overall design accessible, it has been valuable in a working group setting as part of the

continuing design and implementation process.

The analytical part of the approach provides very complete descriptive facilities for the key perspectives of a multi-faceted system design. The service substrate representation makes much of this information accessible, but has not been able to display all of these perspectives. Some hardware-related views, such as network topology or security, are not easily related to the basic services perspective. We have found, however, that having the basic services analysis and protocol details specified, and the service substrate provided as a representation of other system perspectives, makes the entire design accessible in a way that no other approach seems to provide.

### 6.0 Acknowledgements

The techniques described in this paper were evolved and tested during the MIPS system re-design. Much support and encouragement were provided by Bill Green and Tom Handley, as well as the other group leaders. Other members of the MIPS Design Team supplied constructive critiques and suggestions. Joe Kahr, John Duxbury, and other members of the Cassini DCS Working Group provided valuable feedback as the technique was applied to the AMMOS analysis and re-design.

Table 1

MIPS Distributed System Interface	Interface Protocol	Data Units
Specs - PMDS		
Object == R-T Service		
Request	PDF or Icon Chain requesting processing of predicted S/C data	PDF or Icon Chain presented to Exec
Response	Initiate data ingest process using TDS/TOT commands	TDS Private Protocol
Fulfillment	UBR, Catalog updates, R-T Displays, TVSA Images, in appropriate data media or formats	VICAR format files and appropriate service protocols (MDMS, WMS, TVSA)
Object == MRPS Service		
Request	MRPS Request Buffer (File, Processing, Quantity)	MRPS Private (RPC) + VICAR Data File
Response	Accept /ETAY Deny	MRPS Private (RPC)
Fulfillment	Photoproduct written to film	Photoproduct is sent to MPSF for processing, QC'ed and Delivered
Object == MDMS Service		
Request	ANSI-SQL Formatted DBMS query	SQL Std VF protocol or MDMS Gateway protocol
Response	SQL formatted reply/report/result code/Std SQL VF error/status/ data parameters	
Fulfillment		
Object == Mail/SMTP Service		
Request	Process (read/ send/ store/ reply) to RFC822, [MIME] messages	RFC 822, [MIME]
Response	status of process (delivery accomplishment/ failure, [requested reply])	RFC 822 responses
Fulfillment		

Table 2

## Server Interface and Access Specification

		Access Mode											
R.T. Server		Access Mode											
Login FTP File System Process Invocation	Direct Login	Sid VMS	Full User facilities	Ops Users Only *	Sid VMS w/ access control lists	Only Ops users have sid							
	Direct to FTP Daemon	Sid VMS	Full FTP capabilities	Ops users & progs		VMS accts							
	Direct access only	Sid VMS	File system 'unc	Ops users Only		No access from systems outside of MIPS							
	Direct login only	Sid VMS	Full system access	Ops users only		protocol allows execution of limited set of functions							
Systematic Server													
Login FTP File System Process Invocation Processing - remote job request	Direct Login	Sid VMS	Full User facilities	Ops Users Only	Sid VMS w/ access control lists	Only Ops users have sid							
	Direct to FTP Daemon	Sid VMS	Full FTP capabilities	Ops users & progs		VMS accts							
	Direct access only	Sid VMS	File system 'unc	Ops users Only		No access from systems outside of MIPS							
	Direct login	Sid VMS	Full system capabilities	Ops users only		protocol allows execution of limited set of functions							
MIPS Server(s)													
Login FTP File System Film Recorder Processing	Direct Login	Sid VMS	Full User facilities	Ops Users Only	Sid VMS w/ access control lists	Only Ops users have sid							
	Direct to FTP Daemon	Sid VMS	Full FTP capabilities	Any local MIPS users		VMS accts							
	Direct access only	Sid VMS	File system 'unc	Ops users Only	Denial of service if queue fills	Protocol allows execution of limited set of functions							
	MIPS protocol	Kerberos **	Request MIPS services	Any MIPS users									
Database Server													
Login FTP File System DBlib I/F (Sybase) VMS OVS I/F	Direct login/netnet	Sid UNIX/NIS	Full user facilities	Ops Users Only	Sid Unix passwd protection	Only Ops users have sid							
	Direct to FTP Daemon	Sid UNIX/NIS	Full FTP capabilities	Ops users & progs		Unix acct.							
	Direct access only	Sid UNIX/NIS	File system 'unc	Ops users Only		Used by R-T & Syst progs							
	Sybase Protocol	Sybase internal	Full Sybase functionality	Ops users only		Used by all other progs							
Working Mission Storage													
Login FTP File System VMS Archive I/F File Processing	Direct login/netnet	Sid UNIX/NIS	Full user facilities	Ops Users Only	Sid Unix passwd protection	Only Ops users have sid							
	Direct to FTP Daemon	Sid UNIX/NIS	Full FTP capabilities	Ops users & progs		Unix acct.							
	NFS	Sid UNIX/NIS	File system 'unc	Only MIPS internal systems		Used by R-T & Syst progs							
	FEI Protocol	Kerberos	File mode conversion/compression	Any local or remote users		Used by all other progs							
Workstations & Workstations													
Login FTP File System Process Invocation	Direct login/netnet	Sid UNIX/NIS	Full user facilities	Any local MIPS users	Sid Unix passwd protection	Only Ops users have sid							
	Direct to FTP Daemon	Sid UNIX/NIS	Full FTP capabilities	Any local MIPS users		Unix acct.							
	Direct access only	Sid UNIX/NIS	File system 'unc	Any local MIPS users		Used by R-T & Syst progs							
	Direct login	Sid UNIX/NIS	Full system capabilities	Any local MIPS users		Used by all other progs							
PLSS													
Login FTP File System Process Invocation	Direct login/netnet	Sid Unix	Full user facilities	Any local or remote users	Sid Unix passwd protection	Only system open for direct							
	Direct to FTP Daemon	Sid Unix	Full FTP capabilities	Any local or remote users	Corrupt local files and/or system if root	remote user login/access							
	Direct access only	Sid Unix	File system 'unc	Any local or remote users	Remote pwds are sent in the clear on network links	Use local passwd file to isolate remote user accts							
	Direct login	Sid Unix	Full system capabilities	Any local or remote users									

\* Operational users may include analysts, Ops Staff, SAs, or developers, as required

\*\* Implementation details for this function are still TBD, may be handled as part of GUI development

\*\*\* Implementation of Kerberos authentication not yet planned for MRPS